arXiv:1704.00125v1 [cs.DM] 1 Apr 2017

Thin graph classes and polynomial-time approximation schemes

Zdeněk Dvořák*

April 4, 2017

Abstract

Baker [1] devised a powerful technique to obtain approximation schemes for various problems restricted to planar graphs. Her technique can be directly extended to various other graph classes, among the most general ones the graphs avoiding a fixed apex graph as a minor. Further generalizations (e.g., to all proper minor closed graph classes) are known, but they use a combination of techniques and usually focus on somewhat restricted classes of problems. We present a new type of graph decompositions (thin systems of overlays) generalizing Baker's technique and leading to straightforward polynomial-time approximation schemes. We also show that many graph classes (all proper minor-closed classes, and all subgraph-closed classes with bounded maximum degree and strongly sublinear separators) admit such decompositions.

1 Introduction

Baker [1] devised a technique that leads to polynomial-time approximation schemes for a range of problems (including maximum independent set, minimum dominating set, largest H-matching, minimum vertex cover, and many others) when restricted to planar graphs. In modern terms, the technique is based on the fact that if $(V_1, \ldots, V_d)$ is a partition of vertices of a connected planar graph G according to their distance from an arbitrarily chosen vertex, then for any positive integer s and $0 \leq i \leq d - s + 1$, the treewidth of $G[V_{i+1} \cup V_{i+2} \cup \ldots \cup V_{i+s-1}]$ is bounded by a function of s (more specifically, it is at most $3s$), and hence many natural problems can be exactly solved for these subgraphs in linear time [4]. As an example, the exact sizes of largest independent sets in graphs $G - \bigcup_{i:i \bmod s=r} V_i$ for $r = 0, 1, \ldots, s-1$ can be

*Computer Science Institute (CSI) of Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic. E-mail: rakdver@iuuk.mff.cuni.cz. Supported by (FP7/2007-2013)/ERC Consolidator grant LBCAD no. 616787.

found in linear time, and it is easy to see that one of these sets has size at least $(1 - 1/s)$ times the size of the largest independent set in G.

It is natural to ask whether these algorithms can be extended to larger classes $\mathcal{G}$ of graphs. A *layering* of a graph G is a sequence $(V_1, \ldots, V_d)$, where $\{V_1, \ldots, V_d\}$ is a partition of $V(G)$ and G contains no edges between sets V_i and V_j such that $|i-j| \geq 2$. The sets $V_1, \ldots, V_d$ are the *layers* of the layering. Thus, assuming that V_1 intersects every component of G, the most natural layering is obtained by assigning to V_i the vertices at distance exactly $i-1$ from V_1. Let us say that a class $\mathcal{G}$ has *bounded treewidth layerings* if for some function f, each graph $G \in \mathcal{G}$ has a layering in which the union of any s consecutive layers induces a subgraph of treewidth at most $f(s)$. Baker's technique directly extends to all such graph classes (assuming that a suitable layering can be found in polynomial time). A natural obstruction for the existence of bounded treewidth layerings is as follows: let G_n be the graph obtained from the $n \times n$ grid by adding a universal vertex adjacent to all other vertices. Then each layering of G_n has at most three layers, and $\mathrm{tw}(G_n) > n$. Hence, if a class $\mathcal{G}$ has bounded treewidth layerings, then it contains only finitely many of the graphs $\{G_n : n \geq 1\}$. Conversely, Eppstein [14] proved that proper minor-closed classes that avoid some graph G_n have bounded treewidth layerings.

As another obstruction, let U_n denote the graph obtained from an $n \times n \times n$ grid by adding all diagonals to its unit subcubes. Although the graphs U_n have bounded maximum degree and very simple structure, Dvořák et al. [2] proved that for every integer k, there exists n_0 such that for all $n \geq n_0$, the vertex set of U_n cannot be partitioned into two parts both inducing subgraphs of treewidth at most k. This prevents existence of bounded treewidth layerings of U_n (otherwise the partition of the graph into odd and even numbered layers would give a contradiction).

In this paper, we introduce the concept of *thin systems of overlays* which generalizes the notion of bounded treewidth layerings, working around both mentioned obstructions. In particular, we show that all proper minor-closed classes (Corollary 5.16) as well as all subgraph-closed classes with bounded maximum degree and strongly sublinear separators (Corollary 5.22) admit such thin systems of overlays. We also show that thin system of overlays can be used to design simple polynomial-time approximation schemes for many problems: the previously mentioned maximum independent set and minimum dominating set and their distance variants, largest H-matching, and minimum vertex cover, as well as any monotone maximization problems expressible in a restricted fragment of Monadic Second Order logic with constant distance predicates (see Theorem 3.1 for a precise formulation) which includes problems such as maximum c-colorable induced subgraph for any constant number of colors c. Let us remark that in all the cases, the approximation ratio only affects the multiplicative constant of the time complexity of the algorithm (and not the degree of the polynomial), and in

particular the results also imply fixed-parameter tractability of the problems when parameterized by the order of the optimal solution.

The paper is organized as follows: In Section 2 we introduce the notion of thin systems of overlays and study its basic properties. In Section 3 we present the polynomial-time approximation schemes obtained using this notion. Finally, in Section 5, we give a number of operations on graph classes that preserve the existence of thin systems of overlays, and as a corollary we show that all proper minor-closed classes as well as all subgraph-closed classes with bounded maximum degree and strongly sublinear separators admit them.

1.1 Related results

As we already mentioned in the introduction, Eppstein [14] and Demaine and Hajiaghayi [7] showed that Baker's technique generalizes to all proper minor-closed classes that do not contain all apex graphs. Going beyond the apex graph boundary, Grohe [17], Demaine et al. [9] and Dawar et al. [5] generalized the approximation algorithms to all proper minor-closed classes using the tree decomposition from the minor structure theorem. Baker's technique also applies to other geometrically defined graph classes, such as unit disk graphs [20] and graphs embedded with a bounded number of crossings on each edge [16].

Another powerful and largely independent approach to approximation algorithms in proper minor-closed graph classes is through bidimensionality theory, bounding the treewidth of the graph in the terms of the size of the optimal solution and exploiting the arising bounded-size balanced separators to obtain approximate solutions. Demaine and Hajiaghayi [8] and Fomin et al. [15] use this approach to construct polynomial-time approximation schemes on all proper minor-closed classes for many minor-monotone problems (e.g., minimum vertex cover) and on apex-minor-free classes for contraction-monotone problems (e.g., minimum dominating set).

Suppose that $\mathcal{G}$ is a class of graphs with bounded treewidth layerings and k is a positive integer. As a particular corollary of Baker's technique, every graph $G \in \mathcal{G}$ contains pairwise disjoint sets $X_1, \ldots, X_k \subseteq V(G)$ such that removal of any of the sets results in a graph whose treewidth is bounded by a constant (depending on the class of graphs and on k, but not on G itself). In terms introduced by Dvořák [12], such classes are *treewidth-fragile*. As shown by DeVos et al. [10] and algorithmically by Demaine et al. [9], every proper minor-closed class is treewidth-fragile. Treewidth-fragility is by itself sufficient to obtain polynomial-time approximation schemes for some graph parameters such as the independence number or the size of the largest H-matching, but fails for others (minimum dominating set, distance constrained versions of the independence number). Demaine et al. [6] give an edge-contraction variation of this concept, which has similar algorithmic

applications.

Some rather simple graph classes such as 3-dimensional grids with diagonals in unit subcubes are not treewidth-fragile [2]. To overcome this obstruction, Dvořák [12] introduced the fractional version of treewidth-fragility which is still sufficient to obtain the polynomial-time approximation schemes, and showed that all subgraph-closed graph classes with bounded maximum degree and strongly sublinear separators are fractionally treewidth-fragile. The notion of thin systems of overlays can be seen as a further generalization that is able to handle approximation of wider class of problems such as minimum dominating set and distance-r independent set. Let us remark that every class that admits thin systems of overlays is fractionally treewidth-fragile.

A very different approach is taken by Cabello and Gajser [3] for proper minor-closed classes and more generally by Har-Peled and Quanrud [19] for classes of graphs with polynomial expansion (which by the result of Dvořák and Norin [13] is equivalent to having strongly sublinear separators), who showed that trivial local search algorithm gives polynomial-time approximation scheme for problems such as maximum independent and minimum dominating set. It is plausible (but not entirely clear) that this technique can be extended to other first-order definable properties.

In conclusion, our approach seems to be the first one to give a graph decomposition notion generalizing Baker's bounded treewidth layerings that is strong enough to give simple approximation schemes for most problems approachable via Baker's technique, and at the same time can be applied to all proper minor-closed classes as well as many graph classes that are not minor-closed.

Let us remark that the approximation factor does not affect the exponent in the complexity of the obtained approximation schemes, and consequently we also obtain fixed-parameter tractability when parameterized by the order of the optimum solution. However, fixed-parameter tractability of the considered problems is generally known for much larger graph classes, see e.g. [18].

1.2 Limitations of our technique

Most of the algorithmic problems we consider, and in particular maximum independent set and minimum dominating set, are APX-hard even when restricted to graphs with bounded maximum degree [22]. Hence, we cannot hope to generalize our results to any class of graphs that includes say all 3-regular graphs, and in particular we cannot consider topological minor-closed classes instead of minor-closed classes.

All the classes we consider have strongly sublinear separators; and it is easy to see that graphs from every class that admits thin systems of overlays must have sublinear separators. This is rather natural in the context of the

previous paragraph and the theory of bounded expansion: it is known that a subgraph-closed class has polynomial expansion if and only if it has strongly sublinear separators [13], and that subexponential expansion implies existence of sublinear separators [21], while the class of 3-regular graphs has exponential expansion. It would be interesting to see whether this intuition can be made precise and for example show that the maximum independent set problem is APX-hard on any subgraph-closed class of graphs that does not have sublinear separators (or at least has exponential expansion).

We show that subgraph-closed graph classes with strongly sublinear separators admit thin systems of overlays, under the additional assumption that their maximum degree is bounded. It seems plausible that this additional assumption could be dropped, leading to a much more general result strengthening majority of the algorithms mentioned in the previous subsection. Nevertheless, even the simpler question of fractional treewidth-fragility of such graph classes raised in [12] is open.

Finally, let us remark that by its nature as a generalization of Baker's technique, our technique is ill-suited for dealing with global problems, where parts of the solution interact over unlimited distances. For the same reason, the approximation algorithms we design rely on at least a limited form of monotonicity. Thus, we do not obtain approximations for fundamentally non-monotone problems such as minimum independent dominating set or minimum connected dominating set; the bidimensionality theory approach is more suitable in these cases.

2 Thin systems of overlays

Before we give the formal definition of thin systems of overlays, let us start with a motivating example. Given an input graph G, the *minimum r-dominating set* problem (where r is a fixed positive integer) asks for a minimum-size set $X \subseteq V(G)$ such that each vertex of G is at distance at most r from X. Suppose we want to apply Baker's technique to devise a polynomial-time approximation scheme for this problem when restricted to connected *planar* graphs. Let us say that a set $Y \subseteq V(G)$ *r-dominates* a set $Z \subseteq V(G)$ if each vertex of Z is at distance at most r from Y.

Let k be a positive integer. Let us start with a layering $(V_1, \ldots, V_d)$ according to the distance from a vertex of G, and for any integer i, let $G_i = G[V_{i-r} \cup \ldots \cup V_{i+kr+r-1}]$ and $Z_i = V_i \cup \ldots \cup V_{i+kr-1}$ (where $V_j = \emptyset$ if $j \leq 0$ or $j > d$). That is, Z_i consists of kr consecutive layers starting at the i-th one, and G_i is the subgraph of G induced by Z_i and all vertices at distance at most r from the set Z_i. If X is an r-dominating set in G, then the distance between any vertex of Z_i and X is at most r, and thus $X \cap V(G_i)$ r-dominates Z_i in G_i. Conversely, let $Y_i \subseteq V(G_i)$ be a smallest set that r-dominates Z_i in G_i, which can be found in linear time

since G_i has bounded treewidth. Let us fix any $a \in \{0, \ldots, kr-1\}$. Then the sets $\{Z_i : i \bmod kr = a\}$ form a partition of $V(G)$, and thus the set $Y^a = \bigcup_{i \bmod kr = a} Y_i$ is an r-dominating set in G. Since the graphs G_i such that $i \bmod kr = a$ only overlap in r of each rk consecutive layers, it is easy to see that for some $a \in \{0, \ldots, kr-1\}$, the set Y^a is at most $(1 + 1/k)$ times as large as a minimum r-dominating set in G. Thus, the algorithm can just return the smallest of the sets Y_a for $a = 0, \ldots, kr-1$.

This algorithm can be reformulated as follows. For $a = 0, \ldots, kr-1$, let G^a be the graph obtained from graphs G_i such that $i \bmod kr = a$ as their disjoint union (first renaming the vertices that they have in common) and let Z^a be analogously obtained from sets Z_i with $i \bmod kr = a$ (renaming the shared vertices in the same way). Then G^a has bounded treewidth and Y^a can be taken as the smallest set of vertices of G^a that r-dominates Z_a, up to renaming.

Note that due to renaming, G^a is not a subgraph of G; rather, there exists a surjective homomorphism f from G^a to G that describes the renaming (let us recall that a *homomorphism* is a function from $V(G^a)$ to $V(G)$ that maps edges to edges; we only consider simple graphs without loops, and thus images of adjacent vertices must be distinct). For each vertex $x \in Z^a$ and each walk W' of length at most r in G starting in $f(x)$, there exists a walk W of the same length in G^a starting in x such that $f(W) = W'$. This is not the case for vertices in $V(G^a) \setminus Z^a$, for which only smaller neighborhoods are represented in G^a. For each $x \in V(G^a)$, we let $\ell(x)$ denote the size of the represented neighborhood, so we have $\ell(x) = r$ for $x \in V(G^a)$ and $\ell(\cdot)$ gets progressively smaller in the layers away from Z^a. Lastly, it is important that the number of renamed vertices is small compared to $|V(G)|$, and that different vertices get renamed for different values of a.

With these observations in mind, let us proceed with formal definitions.

Definition 1. *Let G and H be graphs, and let $f : V(H) \to V(G)$ and $\ell : V(H) \to \mathbf{Z}_0^+$ be functions. We say that $L = (H, f, \ell)$ is an* overlay *of a graph G if f is a homomorphism from H to G. The overlay is* walk-preserving *if for all $x \in V(H)$ with $\ell(x) \geq 1$ and for each neighbor w of $f(x)$ in G, there exists a neighbor y of x in H such that $f(y) = w$ and $\ell(y) \geq \ell(x) - 1$. For a vertex $v \in V(G)$, the* thickness *$\theta_L(v)$ of the overlay at v is the number $|f^{-1}(v)|$ of vertices of H mapped to v. The* treewidth *of the overlay is equal to the treewidth of H.*

Since f is a homomorphism from H to G, each walk in H corresponds to a walk in G. The converse is not necessarily true, but the property of being walk-preserving implies the converse at least for walks of bounded length.

Lemma 2.1. *Let $L = (H, f, \ell)$ be an overlay of a graph G. If W is a walk in H from a vertex x to a vertex y, then $f(W)$ is a walk in G from $f(x)$ to $f(y)$. Furthermore, if L is walk-preserving and W' is a walk in G from $f(x)$*

of length at most $\ell(x)$, then there exists a walk W in H starting in x such that $f(W) = W'$.

Proof. The first claim is obvious, since f is a homomorphism. For the second claim, let $W' = v_0 v_1 \dots v_t$, where $f(x) = v_0$ and $t \leq \ell(x)$. Let $x_0 = x$. For $i = 1, \dots, t$, the assumption that L is walk-preserving implies that there exists $x_i \in V(H)$ such that $x_{i-1}x_i \in E(H)$, $f(x_i) = v_i$, and $\ell(x_i) \geq \ell(x_{i-1}) - 1 \geq \ell(x) - i$. Then we can set $W = x_0 x_1 \dots x_t$. □

Based on the previous Lemma, the condition from the following definition guarantees that the neighborhood of each vertex up to distance r is represented in the overlay.

Definition 2. *Let r be a positive integer, let G be a graph and let $L = (H, f, \ell)$ be an overlay of G. We say that L is an* r-neighborhood overlay *if L is walk-preserving, $\ell(x) \in \{0, 1, \dots, r\}$ for all $x \in V(H)$, and for all $v \in V(G)$, there exists $x \in f^{-1}(v)$ such that $\ell(x) = r$.*

Let us remark that an r-neighborhood overlay has positive thickness at each vertex. For future use, let us note the following easy fact.

Observation 2.2. *If G' is a subgraph of a graph G and $L = (H, f, \ell)$ is an r-neighborhood overlay of G, then an r-neighborhood overlay L' of G' with at most as large treewidth and thickness can be constructed in time $O(|V(H)| + |E(H)|)$.*

Proof. To construct $L' = (H', f \restriction V(H'), \ell \restriction V(H'))$, let H' be the graph obtained from H by removing all vertices x with $f(x) \in V(G) \setminus V(G')$ and all edges xy with $f(x)f(y) \in E(G) \setminus E(G')$. □

In our motivating example, each component of the graphs G^a was an induced subgraph of G. This is not necessarily the case for general overlays (and imposing this condition would restrict their applicability significantly, as we will see below in Lemma 4.1). Nevertheless, it is convenient to keep track of this special case; we say that an overlay $L = (H, f, \ell)$ of a graph G is *subgraph-based* if the restriction of f to each component H' of H is injective and $f(H')$ is an induced subgraph of G.

Another variation on overlays will be needed to show that clique-sums of graphs from a class admitting thin overlays form another class admitting thin overlays. A vertex of a graph is *simplicial* if its neighborhood induces a clique. Suppose that all vertices of a set $S \subseteq V(G)$ are simplicial in G. We say that an overlay $L = (H, f, \ell)$ of a graph G is *S-simpliciality-preserving* if for every vertex $v \in S$, all vertices of $f^{-1}(v)$ are simplicial in H. The graph $G^\star$ is obtained from G by, for every non-empty-clique K of G, adding a vertex v_K adjacent to vertices in K. Note that if G' is an (induced) subgraph of G, then $(G')^\star$ is an (induced) subgraph of $G^\star$. We define a

$\star$-overlay of G to be a $(V(G^\star) \setminus V(G))$-simpliciality-preserving overlay of $G^\star$.

Let $\mathcal{A}_r(G)$ denote the class of all r-neighborhood overlays of G. Let $\mathcal{S}_r(G)$ denote the class of all subgraph-based r-neighborhood overlays of G. Finally, let $\bigstar_r(G)$ denote the class of all r-neighborhood $\star$-overlays of G.

Definition 3. *Consider any $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. Let r be a positive integer and let G be a graph. A* system of $\mathcal{P}_r$-overlays *of G is a multiset $\mathcal{L}$ of elements of $\mathcal{P}_r(G)$. Let $G' = G^\star$ if $\mathcal{P} = \bigstar$ and $G' = G$ otherwise. The* thickness *$\theta_{\mathcal{L}}(v)$ of a vertex $v \in V(G')$ is $\frac{1}{|\mathcal{L}|}\sum_{L\in\mathcal{L}} \theta_L(v)$, and the thickness $\theta(\mathcal{L})$ of the system $\mathcal{L}$ is the maximum of $\{\theta_{\mathcal{L}}(v) : v \in V(G')\}$. The* treewidth *$tw(\mathcal{L})$ of the system is the maximum of the treewidths of its members.*

We say that the graph G' from the previous definition is the *graph overlaid by $\mathcal{L}$*. The definition of thickness of a system of overlays has a natural probabilistic interpretation which is sometimes convenient: if the system $\mathcal{L}$ has thickness θ and $L = (H, f, \ell)$ is chosen from $\mathcal{L}$ uniformly at random, then for every vertex v, the expected value of $\theta_L(v)$ is at most θ. For approximation algorithms, an important consequence is that for any set $X \subseteq V(G)$, the expected size of $f^{-1}(X)$ is at most $\theta|X|$, where typically we consider the case that θ is close to 1.

As the approximation algorithm inspects all overlays in the system, we also need the system to be small; for any fixed thickness, we need the size of the system for each graph G in the class to be at most polynomial in the number of vertices n of G (actually, for all graph classes we consider, the size will be either constant or polynomial in $e^{(\log\log n)^2}$). Furthermore, we of course need to be able to find such a system of overlays efficiently. With that in mind, we finally introduce the definition of graph classes that admit thin systems of overlays.

Definition 4. *Consider any $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. Let p_2 be a non-decreasing positive polynomial in one variable, let q be a non-decreasing function, and let r, k, and t be positive integers. A class of graphs $\mathcal{G}$* admits (p_2, q)-small $\mathcal{P}_r$-overlays of thickness $1 + 1/k$ and treewidth t *if for every graph $G \in \mathcal{G}$ with n vertices, there exists a system $\mathcal{L}$ of $\mathcal{P}_r$-overlays of G of thickness at most $1 + 1/k$, treewidth at most t, and size at most $p_2(q(n))$.*

Definition 5. *Consider any $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. Let q be a non-decreasing function and let p be a non-decreasing positive polynomial. Let t be a function in two variables. A class of graphs $\mathcal{G}$ is* $(\mathcal{P}, q)$-thin*, with t* bounding the treewidth of its overlays*, if for each positive integers r and k, there exists a non-decreasing positive polynomial p_2 for which $\mathcal{G}$ admits (p_2, q)-small $\mathcal{P}_r$-overlays of thickness $1 + 1/k$ and treewidth $t(r, k)$. We say that $\mathcal{G}$ is* $(\mathcal{P}, p, q)$-thin *if there additionally exists a function f in two variables and an algorithm with time complexity $O(p(n) + n(q(n))^{f(r,k)})$ that for a graph $G \in \mathcal{G}$ with n vertices returns the corresponding system of $\mathcal{P}_r$-overlays.*

We say that the function q is *subpolynomial* if for every $\varepsilon > 0$ and every positive integer m, $(q(n))^m = O(n^\varepsilon)$. We will generally consider only classes that are $(\mathcal{P}, q)$-thin for some subpolynomial q; we will not care about the exact magnitude of the function $f(r,k)$, and we will write the complexity of the algorithm from the definition as $O(p(n) + n\,\mathrm{poly}(q(n)))$, with the understanding that the degree of the polynomial (and the multiplicative constants of the O-notation) may depend on r and k.

Let us remark that in applications, we generally use only the case that $\mathcal{P} = \mathcal{A}$. As already mentioned, the notion of $\bigstar$-thinness is a technicality we introduce to deal with clique-sums (although note also the application in Theorem 3.3); by Observation 2.2, $\star$-overlays can be turned into normal overlays. There is of course a concern that $G^\star$ could be much larger than G, affecting the time complexity of the operations. However, this is not the case by the following lemma, which shows that for thin classes, $G^\star$ is larger only by a constant factor.

Lemma 2.3. *Consider any $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. If a class $\mathcal{G}$ is $(\mathcal{P}, q)$-thin with t bounding the treewidth of its overlays and $c = 4t(1,1)$, then every graph $G \in \mathcal{G}$ has maximum average degree at most c. Consequently, G is c-degenerate, the largest clique in G has size at most $c+1$, and G contains at most $2^c|V(G)|$ non-empty cliques.*

Proof. Consider a graph $G \in \mathcal{G}$ and its subgraph G', and let $\mathcal{L}$ be a system of $\mathcal{A}_1$-overlays of G' of thickness at most 2 and treewidth at most $t(1,1)$, which exists by Definition 5 and Observation 2.2. Let $L = (H, f, \ell)$ be chosen uniformly at random from $\mathcal{L}$. We have

$$\mathrm{E}[|V(H)|] = \mathrm{E}\left[\sum_{v \in V(G')} \theta_L(v)\right] \leq 2|V(G')|$$

by the linearity of expectation. Fix some $L \in \mathcal{L}$ such that $|V(H)| \leq 2|V(G')|$.

Since L is walk-preserving, each vertex $x \in V(H)$ with $\ell(x) > 0$ satisfies $\deg_H(x) \geq \deg_{G'}(f(x))$, and we conclude that

$$\sum_{x \in V(H)} \deg_H(x) \geq \sum_{v \in V(G')} \deg_{G'}(v),$$

and thus $|E(H)| \geq |E(G')|$.

Since H has treewidth at most $t(1,1)$, it is $t(1,1)$-degenerate and its average degree is at most $2t(1,1)$. Consequently, the average degree of G' is

$$\frac{2|E(G')|}{|V(G')|} \leq \frac{4|E(H)|}{|V(H)|} \leq 4t(1,1).$$

Therefore, the maximum average degree of G is at most c, and thus the graph G is c-degenerate. Consider the ordering of vertices of G such that each vertex has at most c neighbors preceding it in the ordering. If K is a clique of G and v is the last vertex of K in the ordering, then all other vertices of K are among the preceding neighbors of v. Hence, $|K| \leq c+1$. Furthermore, each vertex is the last vertex of at most 2^c non-empty cliques, and thus the total number of non-empty cliques of G is at most $2^c|V(G)|$. □

It follows that when we consider polynomial-time algorithms for graphs in a class $\mathcal{G}$ admitting thin overlays of bounded treewidth, it is sufficient to measure their time complexity with respect to the number of vertices, as the number of edges is larger by at most a constant factor.

Let us remark that in the situation of Definition 4, if $|\mathcal{L}| = b$ and $\mathcal{L} = \{L_i : 1 \leq i \leq b\}$ with $L_i = (H_i, f_i, \ell_i)$, then

$$\sum_{i=1}^{b} |V(H_i)| = \sum_{v \in V(G)} \sum_{i=1}^{b} \theta_{L_i}(v) \leq nb\theta(\mathcal{L}) = O(n \operatorname{poly}(q(n))).$$

Furthermore, since all the graphs H_i have bounded treewidth, the number of their edges is linear in the number of their vertices. Hence linear-time manipulations with such systems of overlays can be performed in time $O(n \operatorname{poly}(q(n)))$, and we will neglect mentioning the complexity of such manipulations afterwards. In particular, these facts together with Observation 2.2 imply that every $(\bigstar, p, q)$-thin class is $(\mathcal{A}, p, q)$-thin.

3 Approximation algorithms via thinness

Let us now show how thin systems of overlays can be used to design approximation schemes. Firstly, we will consider problems such as the following ones:

- Maximum r-independent set, for a fixed positive integer r: Given a graph G, find a largest subset X of vertices of G such that the distance between distinct elements of X is at least r.
- Maximum F-matching, for a fixed connected graph F: Given a graph G, find a largest subset X of vertices of G such that $G[X]$ can be partitioned into vertex-disjoint copies of F.
- Maximum c-colorable induced subgraph, for a fixed positive integer c: Given a graph G, find a largest subset X of vertices of G such that the induced subgraph $G[X]$ can be partitioned into c independent sets.
- The r-distance variants of the previous two problems, for any fixed integer r.

All these problems can be expressed in uniform way as follows. For a positive integer r, let MSO_r denote Monadic Second Order logic with predicates $d_l(u,v)$ for $l = 1, \ldots, r$, where $d_l(u,v)$ is interpreted to be true if there exists a path of length at most l between vertices u and v. A formula φ in MSO_r with one free vertex set variable X is *solution-restricted* if all its quantifications are restricted to X, i.e., φ only contains quantifiers of form $(\exists/\forall v : v \in X)$, $(\exists/\forall V : V \subseteq X)$, $(\exists/\forall e = uv : u, v \in X)$, and $(\exists/\forall E : u, v \in X \text{ for all } uv \in E)$.

Consider the problem of finding a largest set $X \subseteq V(G)$ in a graph G satisfying some property π. For a non-negative integer s, we say that the problem is *s-near-monotone solution-restricted MSO_r-expressible* if the following conditions are satisfied.

- For every set $X \subseteq V(G)$ that has the property π and every set $Y \subseteq X$, there exists a subset of Y of size at least $|Y| - s|X \setminus Y|$ that has the property π.
- There exists a solution-restricted formula φ in MSO_r with one free vertex set variable X such that $G \models \varphi(X)$ if and only if X has the property π.

As an example, consider the r-distance variant of maximum F-matching, for a positive integer r and a connected graph F. Here, π is the property that there exists a set of edges $Z \subseteq E(G[X])$ such that the graph (X, Z) is a disjoint union of copies of F and no two vertices belonging to different components of the graph have distance less than r in G; this property is clearly solution-restricted MSO_r-expressible. Furthermore, if X has this property π and $Y \subseteq X$, then let Y' be the subset of Y obtained by removing all components of the graph (X, Z) that intersect $X \setminus Y$. Clearly, Y' has property π and $|Y'| \geq |Y| - (|V(F)| - 1)|X \setminus Y|$, and thus π is $(|V(F)| - 1)$-near-monotone.

Theorem 3.1. *For integers $s \geq 0$ and $r \geq 1$, let π be an s-near-monotone solution-restricted MSO_r-expressible property, and let $\varepsilon > 0$ be any real number. If a class of graphs $\mathcal{G}$ is $(\mathcal{A}, p, q)$-thin for some functions p and q, then given an n-vertex graph $G \in \mathcal{G}$, a set $X \subseteq V(G)$ satisfying the property π whose size is within $1 - \varepsilon$ factor of the maximum one can be found in time $O(p(n) + n\, poly(q(n)))$.*

Proof. Let φ be a solution-restricted MSO_r-formula expressing the property π. Let k be a positive integer such that $1/k \leq \varepsilon/(s+1)$. Let $\mathcal{L}$ be a system of $\mathcal{A}_r$-overlays of G of thickness at most $1 + 1/k$, treewidth at most $t = t(r,k)$ and size $\mathrm{poly}(q(n))$, which can be found in time $O(p(n) + n\,\mathrm{poly}(q(n)))$.

Consider any overlay $L = (H, f, \ell) \in \mathcal{L}$, and let $S_L = \{v \in V(G) : \theta_L(v) = 1\}$. Note that $\ell(x) = r$ for all $x \in f^{-1}(S_L)$, and that f is an isomorphism from $H[f^{-1}(S_L)]$ to $G[S_L]$. Since φ is solution-restricted, only

the distance predicates in φ take into account the parts of the graph not contained in the considered set; hence, Lemma 2.1 implies that for any $Z \subseteq S_L$, we have $G \models \varphi(Z)$ if and only if $H \models \varphi(f^{-1}(Z))$. Since H has bounded treewidth, we can in time $O(|V(H)|)$ find a set $Z_L \subseteq S_L$ of largest size such that $H \models \varphi(f^{-1}(Z_L))$, and thus Z_L has property π. The algorithm returns the largest set Z_L over all $L \in \mathcal{L}$.

Consider an optimal solution, i.e., a set $X_0 \subseteq V(G)$ of maximum size satisfying the property π. For any $L = (H, f, \ell) \in \mathcal{L}$, let $Y_L = X_0 \cap S_L$. Since π is s-near-monotone, there exists a set $Z'_L \subseteq Y_L$ of size at least $|Y_L| - s|X_0 \setminus Y_L|$ satisfying property π; since $Z'_L \subseteq S_L$, the choice of Z_L implies

$$|Z_L| \geq |Z'_L| \geq |X_0| - (s+1)|X_0 \setminus Y_L| = |X_0| - (s+1)|X_0 \setminus S_L|.$$

Consider any vertex $v \in V(G)$. Since $\theta_{\mathcal{L}}(v) \leq 1 + 1/k$, and $\theta_L(v) \geq 1$ for all overlays $L \in \mathcal{L}$, and $\theta_L(v) \geq 2$ if $v \notin S_L$, we conclude that there are at most $|\mathcal{L}|/k$ overlays $L \in \mathcal{L}$ such that $v \notin S_L$. Consequently,

$$\frac{1}{|\mathcal{L}|} \sum_{L \in \mathcal{L}} |X_0 \setminus S_L| = \sum_{v \in X_0} \frac{1}{|\mathcal{L}|} |\{L \in \mathcal{L} : v \notin S_L\}| \leq \frac{|X_0|}{k},$$

and thus there exists an overlay $L \in \mathcal{L}$ such that $|X_0 \setminus S_L| \leq |X_0|/k$. For this overlay, we have

$$|Z_L| \geq |X_0| - \frac{s+1}{k}|X_0| \geq (1-\varepsilon)|X_0|,$$

and thus the algorithm returns a set whose size is within $1 - \varepsilon$ factor of the maximum one. □

A similar argument can be applied to deal with domination problems.

Theorem 3.2. *Let r be a positive integer and ε a positive real number. If a class of graphs $\mathcal{G}$ is $(\mathcal{A}, p, q)$-thin for some functions p and q, then given an n-vertex graph $G \in \mathcal{G}$, an r-dominating set $X \subseteq V(G)$ whose size is within $1+\varepsilon$ factor of the minimum one can be found in time $O(p(n)+n\,\mathrm{poly}(q(n)))$.*

Proof. Let k be a positive integer such that $1/k \leq \varepsilon$. Let $\mathcal{L}$ be a system of $\mathcal{A}_r$-overlays of G of thickness at most $1 + 1/k$, treewidth at most $t = t(r, k)$ and size $\mathrm{poly}(q(n))$, which can be found in time $O(p(n) + n\,\mathrm{poly}(q(n)))$.

Consider any overlay $L = (H, f, \ell) \in \mathcal{L}$, and let $S_L = \{x \in V(H) : \ell(x) = r\}$. Since L is an r-neighborhood overlay, we have $f(S_L) = V(G)$. Let Z_L be a minimum subset of $V(H)$ that r-dominates S_L in H, which can be found in time $O(|V(H)|)$ since H has bounded treewidth. By Lemma 2.1, $f(Z_L)$ is an r-dominating set in G. The algorithm returns the smallest set $f(Z_L)$ over all $L \in \mathcal{L}$.

Let $X_0 \subseteq V(G)$ be a minimum r-dominating set of G. For an overlay $L = (H, f, \ell) \in \mathcal{L}$, let $Z'_L = f^{-1}(X_0)$. By Lemma 2.1, the set Z'_L r-dominates S_L in H, and thus $|Z'_L| \geq |Z_L|$. Furthermore, $|Z'_L| = \sum_{v \in X_0} \theta_L(v)$. Since $\mathcal{L}$ has thickness at most $1 + 1/k$, there exists an overlay $L \in \mathcal{L}$ such that $|Z'_L| \leq (1 + 1/k)|X_0| \leq (1 + \varepsilon)|X_0|$. Consequently, $|f(Z_L)| \leq |Z_L| \leq |Z'_L| \leq (1 + \varepsilon)|X_0|$, and thus the algorithm returns a set whose size is within $1 + \varepsilon$ factor of the maximum one. □

The notion of $\bigstar$-thinness is mostly a technical tool to deal with clique-sums. Nevertheless, it also gives us a control over the cliques in the overlaid graph, which is useful in the problem of finding a minimum vertex cover, or more generally in the following problem. An *s-clique cover* is a set of vertices of G that intersects all cliques of size at least s.

Theorem 3.3. *Let s be a positive integer and ε a positive real number. If a class of graphs $\mathcal{G}$ is $(\bigstar, p, q)$-thin for some functions p and q, then given an n-vertex graph $G \in \mathcal{G}$, an s-clique cover whose size is within $1 + \varepsilon$ factor of the minimum one can be found in time $O(p(n) + n\, poly(q(n)))$.*

Proof. Let S be the set of vertices of $G^\star$ representing the cliques of G of size s. Let k be a positive integer such that $1/k \leq \varepsilon$. Let $\mathcal{L}$ be a system of $\bigstar_1$-overlays of G of thickness at most $1 + 1/k$, treewidth at most $t = t(1, k)$ and size $\mathrm{poly}(q(n))$, which can be found in time $O(p(n) + n\,\mathrm{poly}(q(n)))$.

Consider any overlay $L = (H, f, \ell) \in \mathcal{L}$. Let $S_L = \{x \in f^{-1}(S) : \ell(x) = 1\}$. Since L is S-simpliciality-preserving, the neighborhood of each vertex $x \in S_L$ induces a clique, and since $\ell(x) = 1$ and L is walk-preserving, the size of this clique is exactly s. Let Z_L be a minimum subset of $V(H)$ intersecting the neighborhoods of all vertices of S_L, which can be found in time $O(|V(H)|)$ since H has bounded treewidth. Since L is a $\star$-overlay of G, $f(Z_L)$ is an s-clique cover of G. The algorithm returns the smallest set $f(Z_L)$ over all $L \in \mathcal{L}$.

Let $X_0 \subseteq V(G)$ be a minimum s-clique cover of G. For any overlay $L = (H, f, \ell) \in \mathcal{L}$, let $Z'_L = f^{-1}(X_0)$, and observe that Z'_L intersects the neighborhoods of all vertices of S_L. Hence, $|Z'_L| \geq |Z_L|$. Furthermore, $|Z'_L| = \sum_{v \in X_0} \theta_L(v)$. Since $\mathcal{L}$ has thickness at most $1 + 1/k$, there exists an overlay $L \in \mathcal{L}$ such that $|Z'_L| \leq (1 + 1/k)|X_0| \leq (1 + \varepsilon)|X_0|$. Consequently, $|f(Z_L)| \leq |Z_L| \leq |Z'_L| \leq (1 + \varepsilon)|X_0|$, and thus the algorithm returns a set whose size is within $1 + \varepsilon$ factor of the maximum one. □

Regarding the time complexity of the algorithms, for all the classes of graphs that we consider the function q is subpolynomial, and thus the time complexity turns out to be at most $O(p(n) + n^{1+\delta})$ for any fixed $\delta > 0$, with only the multiplicative constant depending on the desired precision of approximation.

4 $\mathcal{S}$-thinness and $\bigstar$-thinness

As we have seen at the beginning of Section 2, subgraph-based overlays are an interesting special case arising in previously studied settings, and in many ways $\mathcal{S}$-thinness is easier to work with and might give stronger algorithmic consequences. Nevertheless, as we will show momentarily, only classes with locally bounded treewidth can be $\mathcal{S}$-thin (recall that a class $\mathcal{G}$ has *locally bounded treewidth* if for each positive integer r, the treewidth of r-neighborhoods of vertices in graphs from $\mathcal{G}$ is bounded by a constant depending only on r and the class). This restricts the applicability of the concept, and in particular shows that not all proper minor-closed classes are $\mathcal{S}$-thin.

Lemma 4.1. *For any function q, if a class of graphs $\mathcal{G}$ is $(\mathcal{S}, q)$-thin, then $\mathcal{G}$ has locally bounded treewidth.*

Proof. Let t be the function bounding the treewidth of overlays of $\mathcal{G}$. Consider any positive integer r. Every graph $G \in \mathcal{G}$ has a system $\mathcal{L}$ of $\mathcal{S}_r$ overlays of thickness at most 2 and treewidth at most $t(r, 1)$. Consider any overlay $L = (H, f, \ell) \in \mathcal{L}$ and any vertex $v \in V(G)$. Since L is an r-neighborhood overlay of G, there exists a vertex $x \in V(H)$ such that $f(x) = v$ and $\ell(x) = r$. Since L is subgraph-based, Lemma 2.1 implies that the component H' of H with $x \in V(H')$ has a subgraph isomorphic to the subgraph of G induced by vertices at distance at most r from v. Hence, this subgraph must have treewidth at most $\mathrm{tw}(H') \leq t(r, 1)$. As the choice of G, v and r was arbitrary, we conclude that $\mathcal{G}$ has locally bounded treewidth. □

Trivially, $\mathcal{S}$-thinness implies $\mathcal{A}$-thinness. Let us also remark that $\mathcal{S}$-thinness implies $\bigstar$-thinness.

Lemma 4.2. *For any functions p and q, if a class of graphs $\mathcal{G}$ is $(\mathcal{S}, p, q)$-thin, then $\mathcal{G}$ is also $(\bigstar, p, q)$-thin.*

Proof. Let the function t bound the treewidth of overlays of $\mathcal{G}$. Consider any positive integers k and r, and a graph $G \in \mathcal{G}$ with n vertices. Let $\mathcal{L}$ be the system of $\mathcal{S}_r$-overlays of G of thickness at most $1 + 1/k$, found in time $O(p(n) + n\,\mathrm{poly}(q(n)))$ according to Definition 5. Consider any overlay $L = (H, f, \ell) \in \mathcal{L}$. Let $L_1 = (H_1, f_1, \ell_1)$ be obtained from L as follows. The graph H_1 is a supergraph of H with functions f_1 and ℓ_1 matching f and ℓ on $V(H)$. For each non-empty clique K of G and each component H' of H such that $K \subseteq f(V(H'))$, we add a new vertex x to H_1 with neighborhood $V(H') \cap f^{-1}(K)$. We set $f_1(x) = v_K$, where v_K is the vertex of $G^\star$ representing the clique K. Also, we set $\ell_1(x) = \max\{\ell(y) : y \in V(H') \cap f^{-1}(K)\}$.

Let us verify that L_1 is an r-neighborhood $\star$-overlay of G. By construction, it is a $(V(G^\star) \setminus V(G))$-simpliciality-preserving overlay of $G^\star$. Let x

be any vertex of H_1 with $\ell_1(x) \geq 1$ and let w be a neighbor of $f_1(x)$ in $G^\star$. Note that at most one of $f_1(x)$ and w belongs to $V(G^\star) \setminus V(G)$. If $f_1(x), w \in V(G)$, then x has a neighbor y such that $f_1(y) = f(y) = w$ and $\ell_1(y) = \ell(y) \geq \ell(x) - 1 = \ell_1(x) - 1$ since L is a walk-preserving overlay of G. If $w \in V(G^\star) \setminus V(G)$, then let H' be the component of H containing x and let K be the clique in G such that $w = v_K$. Note that $f_1(x) = f(x) \in K$, and since $\ell(x) = \ell_1(x) \geq 1$ and L is a subgraph-based overlay, we conclude that $K \subseteq f(V(H'))$; hence, a neighbor y of x such that $f_1(y) = w$ and $\ell_1(y) \geq \ell_1(x)$ was added to H_1 during its construction. If $f_1(x) \in V(G^\star) \setminus V(G)$, then let K be the clique in G such that $f_1(x) = v_K$; we have $w \in K$, and by the construction of H_1, x has a unique neighbor y in H_1 such that $f_1(y) = w$. Let H' be the component of H containing y. Since L is a subgraph-based walk-preserving overlay, note that $\ell(f^{-1}(K) \cap V(H'))$ consists of at most two consecutive integers, and thus $\ell_1(y) \geq \ell_1(x) - 1$. Consequently, the overlay L_1 is walk-preserving. Consider any non-empty clique K of G and a vertex $z \in K$. Since L is an r-neighborhood overlay, there exists a vertex $x \in V(H)$ with $f(x) = z$ and $\ell(x) = r$. Let H' be the component of H containing x. Since L is subgraph-based and walk-preserving, we conclude that $K \subseteq f(V(H'))$, and thus during the construction of H_1, we added a neighbor w of x such that $f_1(w) = v_K$ and $\ell_1(w) = r$. Therefore, L_1 is an r-neighborhood overlay.

Consequently, $L_1 \in \bigstar_r(G)$. Note that L_1 has treewidth at most $\mathrm{tw}(L) + 1$. Furthermore, the system $\{L_1 : L \in \mathcal{L}\}$ has thickness at most $1 + 1/k$, since the thickness at each vertex v_K is bounded by the thicknesses of the vertices of K. Since the described transformation can be performed in time $O(n \operatorname{poly}(q(n)))$, we conclude that $\mathcal{G}$ is $(\bigstar, p, q)$-thin. □

5 Thinness of graph classes

Finally, we show that many important graph classes are thin. Let us start by establishing some basic tools.

5.1 Transformations of thin classes

When considering systems of overlays for several graphs, it is sometimes convenient to have all of them of the same size, which is possible by the following lemma.

Lemma 5.1. *Consider any $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. Let $\mathcal{G}$ be a $(\mathcal{P}, p, q)$-thin class of graphs for some functions p and q, with t bounding the treewidth of its overlays. Let $G_1, \ldots, G_m$ be graphs from $\mathcal{G}$, each with at most n vertices, and let $N = \sum_{i=1}^{m} |V(G_i)|$. For any positive integers r and k, we can find systems $\mathcal{L}_1, \ldots, \mathcal{L}_m$ of $\mathcal{P}_r$-overlays of graphs $G_1, \ldots, G_m$ of thickness at*

most $1 + 1/k$ *and treewidth at most* $t(r, 3k)$ *such that* $|\mathcal{L}_1| = \ldots = |\mathcal{L}_m| \leq$ *poly*$(q(n))$ *in time* $O(p(N) + N$ *poly*$(q(n)))$.

Proof. For $1 \leq i \leq m$, since the class $\mathcal{G}$ is $(\mathcal{P}, p, q)$-thin, we can find a system $\mathcal{L}'_i$ of $\mathcal{P}_r$-overlays of G_i of thickness at most $1+\frac{1}{3k}$, treewidth at most $t(r, 3k)$, and order at most poly$(q(n))$ in time $O(p(|V(G_i)|)+|V(G_i)|\text{poly}(q(n)))$. Let $a = \max\{|\mathcal{L}_i| : 1 \leq i \leq m\}$. For $i \in \{1, \ldots, m\}$, let $\mathcal{L}_i$ be a multiset obtained from $\mathcal{L}'_i$ by replicating its elements so that $|\mathcal{L}_i| = 3ka \leq \text{poly}(q(n))$ and the number of times an element of $\mathcal{L}'_i$ is replicated is either c_i or $c_i + 1$ for some positive integer c_i. The total time needed to obtain these systems clearly is $O(p(N)+N\,\text{poly}(q(n)))$, hence we just need to show that they have thickness at most $1 + 1/k$.

Consider any $i \in \{1, \ldots, m\}$, and let $\mathcal{L}'_i = \{L_1, \ldots, L_{b_i}\}$, where $b_i \leq a$. Note that $3ka = |\mathcal{L}_i| \geq b_i c_i$. If v is any vertex of the graph overlaid by $\mathcal{L}_i$, then

$$\begin{aligned}
\theta_{\mathcal{L}_i}(v) &\leq \frac{c_i + 1}{3ka} \sum_{j=1}^{b_i} \theta_{L_j}(v) \\
&= \frac{b_i(c_i + 1)}{3ka} \theta_{\mathcal{L}'_i}(v) \\
&\leq \Big(1 + \frac{b_i}{3ka}\Big)\theta_{\mathcal{L}'_i}(v) \leq \Big(1 + \frac{1}{3k}\Big)\theta_{\mathcal{L}'_i}(v) \\
&\leq \Big(1 + \frac{1}{3k}\Big)^2 < 1 + \frac{1}{k},
\end{aligned}$$

as required. □

For $1 \leq i \leq m$, let $L_i = (H_i, f_i, \ell_i)$ be an overlay of a graph G. We define $L_1 \circ \ldots \circ L_m$ as the overlay (H, f, ℓ) of G, where H is the disjoint union of $H_1, \ldots, H_m$, and for each $i \in \{1, \ldots, m\}$ and each vertex $v \in V(G_i)$ we have $f(v) = f_i(v)$ and $\ell(v) = \ell_i(v)$.

Let $\mathcal{L}_1, \ldots, \mathcal{L}_m$ be systems of overlays of a graph G such that $|\mathcal{L}_1| = \ldots = |\mathcal{L}_m| = a$. For $i \in \{1, \ldots, m\}$, let $\mathcal{L}_i = \{L_i^1, \ldots, L_i^a\}$, with the elements listed in an arbitrary order. For $j \in \{1, \ldots, a\}$, let $L^j = L_1^j \circ \ldots \circ L_m^j$, and we define $\mathcal{L}_1 \circ \ldots \circ \mathcal{L}_m$ to denote the system of overlays $\{L^1, \ldots, L^a\}$. We also use $\bigcirc_{i=1}^m \mathcal{L}_i$ to denote $\mathcal{L}_1 \circ \ldots \circ \mathcal{L}_m$ and other analogous $\sum$-like notation for the $\circ$ operation.

Lemma 5.2. *Consider any* $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. *Let* $\mathcal{G}$ *be a* $(\mathcal{P}, p, q)$*-thin class of graphs for some functions* p *and* q*, and let* $\mathcal{G}'$ *be a class of graphs. If for every* $G \in \mathcal{G}'$*, each component of* G *belongs to* $\mathcal{G}$*, then* $\mathcal{G}'$ *is* $(\mathcal{P}, p, q)$*-thin.*

Proof. Let t bound the treewidth of the overlays of $\mathcal{G}$. Let r and k be positive integers. Let an n-vertex graph $G \in \mathcal{G}'$ have components $G_1, \ldots, G_m \in \mathcal{G}$, and let $\mathcal{L}_1, \ldots, \mathcal{L}_m$ be the systems of $\mathcal{P}_r$-overlays of $G_1, \ldots, G_m$ of thickness

at most $1+1/k$ and treewidth at most $t(r,3k)$, such that $|\mathcal{L}_1| = \ldots = |\mathcal{L}_m| \leq \mathrm{poly}(q(n))$, found in total time $O(p(n)+n\,\mathrm{poly}(q(n)))$ by Lemma 5.1. Then $\mathcal{L} = \mathcal{L}_1 \circ \ldots \circ \mathcal{L}_m$ is a system of $\mathcal{P}_r$-overlays of G of thickness at most $1+1/k$ and treewidth at most $t(r,3k)$ with $|\mathcal{L}| = |\mathcal{L}_1| \leq \mathrm{poly}(q(n))$. □

We will need the following fact.

Observation 5.3. *Let x, y, m_x, and m_y be integers such that $|m_x - m_y| \leq 1$. If $y \geq x-1$, then $\min(y, m_y) \geq \min(x, m_x) - 1$.*

We now consider graphs with layerings such that unions of any few layers belong to a thin class.

Lemma 5.4. *Consider any $\mathcal{P} \in \{\mathcal{A}, \mathcal{S}, \bigstar\}$. For every positive integer i, let $\mathcal{G}_i$ be a $(\mathcal{P},p,q)$-thin class of graphs. Let $\mathcal{G}$ be a class of graphs such that every $G \in \mathcal{G}$ has a layering $(V_1,\ldots,V_d)$ such that for $1 \leq i \leq d$, the union of each at most i consecutive layers induces a subgraph of G belonging to $\mathcal{G}_i$. Suppose furthermore that such a layering can be found in time $O(p(n)+n\,poly(q(n))$, where $n = |V(G)|$. Then $\mathcal{G}$ is a $(\mathcal{P},p,q)$-thin.*

Proof. For a positive integer i, let t_i denote the function bounding the treewidth of overlays of $\mathcal{G}_i$. Let $G \in \mathcal{G}$ be an n-vertex graph, let r and k be positive integers, and let $(V_1,\ldots,V_d)$ be the layering of G from the assumptions of the lemma. For any integer i such that $i \leq 0$ or $i > d$, let $V_i = \emptyset$.

We need to find a system of $\mathcal{P}_r$-overlays of G of thickness at most $1+\frac{1}{k}$, size at most $\mathrm{poly}(q(n))$, and bounded treewidth, in time $O(p(n)+n\,\mathrm{poly}(q(n))$. Let $\Delta = 6kr$. If $d \leq \Delta$, then $G \in \mathcal{G}_\Delta$ and such a system can be found using the algorithm for the class G_Δ. Hence, assume that $d \geq \Delta$.

For an integer j, let G_j be the subgraph of G induced by $V_{j-r} \cup \ldots \cup V_{j+\Delta+r-1}$. Note that $G_j \in \mathcal{G}_{\Delta+2r}$. Using Lemma 5.1, for each integer j such that $V(G_j) \neq \emptyset$ we obtain a system $\mathcal{L}_j$ of $\mathcal{P}_r$-overlays of G_j of thickness at most $1+\frac{1}{2k}$ and treewidth at most $t_{\Delta+2r}(r,6k)$, such that for all such integers j the size of the system $\mathcal{L}_j$ is the same (and bounded by $\mathrm{poly}(q(n))$). Since each vertex of G belongs to at most $\Delta+2r$ of the graphs G_j, all these systems of overlays can be constructed in time $O(p(n)+n\,\mathrm{poly}(q(n)))$.

Consider an integer j with $V(G_j) \neq \emptyset$, and an overlay $L = (H,f,\ell) \in \mathcal{P}_r(G_j)$. For $x \in V(H)$ such that $f(x) \in V_i$ for some integer i, let $m_L(x) = r$ if $j \leq i \leq j+\Delta-1$, $m_L(x) = r-(j-i)$ if $i < j$ and $m_L(x) = r-(i-(j+\Delta-1))$ if $i > j+\Delta-1$; note that $m(x) \geq 0$. In the case that $\mathcal{P} = \bigstar$, for $x \in V(H)$ such that $f(x) \in V(G_j^\star) \setminus V(G_j)$, let $m_L(x)$ be the maximum of $\{m_L(y) : xy \in E(H)\}$; note that the neighborhood $N(x)$ of x in H induces a clique, and thus $f(N(x))$ is contained in at most two consecutive layers of the layering, and hence $m_L(N(x))$ consists of at most two consecutive integers. Let the system $\mathcal{L}'_j$ be obtained from $\mathcal{L}_j$ as follows: for each $L = (H,f,\ell) \in \mathcal{L}_j$, let us replace ℓ by the function ℓ' defined by $\ell'(x) = \min(\ell(x), m_L(x))$

for all $x \in V(H)$. Let $L' = (H, f, \ell')$. Since $\ell'(x) = 0$ for all x such that $f(x) \in V_{j-r} \cup V_{j+\Delta-1+r}$ and $(V_1, \ldots, V_d)$ is a layering, Observation 5.3 implies that L' taken as an overlay of G (or $G^\star$) and not just of G_j (or $G_j^\star$) is walk-preserving.

Let us consider an integer b such that $0 \leq b \leq \Delta - 1$, and for any integer j such that $j \bmod \Delta = b$ and $V(G_j) \neq \emptyset$, let $L'_j = (H_j, f_j, \ell'_j)$ be any element of $\mathcal{L}'_j$, obtained according to the previous paragraph from $L_j = (H_j, f_j, \ell_j)$. Let

$$L^\circ = \bigcirc_{\substack{j \bmod \Delta = b, \\ V(G_j) \neq \emptyset}} L'_j.$$

We claim that $L^\circ = (H^\circ, f^\circ, \ell^\circ)$ belongs to $\mathcal{P}_r(G)$. Indeed, consider a vertex v of the graph overlaid by L°; we need to show that there exists a vertex $x \in V(H^\circ)$ such that $f^\circ(x) = v$ and $\ell^\circ(x) = r$. If $v \in V(G)$, then let $v' = v$; if $\mathcal{P} = \bigstar$ and $v \in V(G^\star) \setminus V(G)$, then let v' be an arbitrary neighbor of v in $V(G)$, contained in the clique represented by v. Let V_i be the layer containing v'. There exists unique j such that $j \leq i \leq j+\Delta-1$ and $j \bmod \Delta = b$. Note that $v \in V(G_j^\star)$. Since L_j belongs to $\mathcal{P}_r(G_j)$, there exists $x \in V(H_j)$ such that $f_j(x) = v$ and $\ell_j(x) = r$. By the choice of ℓ'_j, observe that $\ell^\circ(x) = \ell'_j(x) = \ell_j(x) = r$. Hence, L° is an r-neighborhood overlay of G if $\mathcal{P} \neq \bigstar$ and of $G^\star$ if $\mathcal{P} = \bigstar$. In the case that $\mathcal{P} = \bigstar$, furthermore note that since each overlay L'_j is $(V(G_j^\star) \setminus V(G_j))$-simpliciality-preserving, L° is $(V(G^\star) \setminus V(G))$-simpliciality-preserving.

For $0 \leq b \leq \Delta - 1$, let

$$\mathcal{L}^b = \bigcirc_{\substack{j \bmod \Delta = b, \\ V(G_j) \neq \emptyset}} \mathcal{L}'_j.$$

By the previous paragraph, $\mathcal{L}^b$ is a system of $\mathcal{P}_r$-overlays for G, and observe that the treewidth of $\mathcal{L}^b$ is at most $t_{\Delta+2r}(r, 6k)$. Let $D_b = \{l : l \equiv b - r, b - r + 1, \ldots, b + r - 1 \pmod{\Delta}\}$. Observe that for any index $l \in D_b$ and a vertex $v \in V_l$, there exist two indices j such that $j \bmod \Delta = b$ and $v \in V(G_j)$, while for $l \notin D_b$ and $v \in V_l$, there exists only one such index. Hence, $\theta_{\mathcal{L}^b}(v) \leq 2 + \frac{1}{k} \leq 1 + \frac{1}{2k} + 3/2$ in the former case and $\theta_{\mathcal{L}^b}(v) \leq 1 + \frac{1}{2k}$ in the latter case. Similarly, if $\mathcal{P} = \bigstar$, v is a vertex of $V(G^\star) \setminus V(G)$, and v has a neighbor in V_l for some $l \notin D_b$, then $v \in V(G_j^\star)$ for a unique index j such that $j \bmod \Delta = b$, and otherwise $v \in V(G_j^\star)$ for two such indices j, with the same resulting bounds on the thickness of v.

Let $\mathcal{L} = \bigcup_{b=0}^{\Delta-1} \mathcal{L}^b$. Note that each integer l belongs to D_b for exactly $2r$ values of b, and thus for each vertex v of the graph overlaid by $\mathcal{L}$, we have

$$\theta_{\mathcal{L}}(v) = 1 + \frac{1}{2k} + \frac{3}{2} \cdot \frac{2r}{\Delta} = 1 + 1/k.$$

Therefore, $\mathcal{L}$ is a system of $\mathcal{P}_r$-overlays of G of thickness at most $1 + \frac{1}{k}$, size at most $\mathrm{poly}(q(n))$, and treewidth at most $t_{\Delta+2r}(r, 6k)$, as required. □

The next lemma deals with apex vertices; as such, it cannot apply to subgraph-based overlays by Lemma 4.1, since adding a universal vertex makes locally bounded treewidth impossible (unless the class has bounded treewidth).

Lemma 5.5. *Consider any $\mathcal{P} \in \{\mathcal{A}, \bigstar\}$. Let $\mathcal{G}$ be a $(\mathcal{P}, p, q)$-thin class of graphs for some functions p and q and let a be a positive integer. Let $\mathcal{G}'$ be a class of graphs. Suppose that for each $G \in \mathcal{G}'$, there exists a set $A \subseteq V(G)$ of size at most a such that $G - A \in \mathcal{G}$. Suppose furthermore that such a set A can be found in time $O(p(n) + n\,poly(q(n)))$, where $n = |V(G)|$. Then $\mathcal{G}'$ is $(\mathcal{P}, p, q)$-thin.*

Proof. Let t bound the treewidth of overlays of $\mathcal{G}$. Let G be an n-vertex graph from $\mathcal{G}'$, let $A \subseteq V(G)$ be a set of size at most a such that $G - A \in \mathcal{G}$, and let r and k be positive integers. Let $\mathcal{L}$ be a system of $\mathcal{P}_r$-overlays of $G - A$ of thickness at most $1 + 1/k$, treewidth at most $t(r, k)$, and size at most $\mathrm{poly}(q(n))$, found in time $O(n\,\mathrm{poly}(q(n)))$. Let $A = \{v_1, \ldots, v_b\}$ for some $b \leq a$.

Consider any overlay $L = (H, f, \ell) \in \mathcal{L}$. Let H', f', and ℓ' be obtained from H, f, and ℓ as follows: to H, we add vertices x_1, ..., x_b with $\ell'(x_1) = \ldots = \ell'(x_b) = r$, edges between vertices x_i and x_j for all $i, j \in \{1, \ldots, b\}$ such that $v_i v_j \in E(G)$, and edges $x_i y$ for all $i \in \{1, \ldots, b\}$ and $y \in V(H)$ such that $v_i f(y) \in E(G)$, and we set $f'(x_i) = v_i$ for $1 \leq i \leq b$. Additionally, if $\mathcal{P} = \bigstar$, we modify the resulting overlay as follows. For each non-empty clique K of $G[A]$, we add a vertex x with $\ell'(x) = r$ adjacent to vertices of $(f')^{-1}(K)$ and let $f'(x_K)$ be the corresponding vertex of $G^\star$. Also, for each clique K of G such that $K \cap A \neq \emptyset$ and $K \not\subseteq A$, letting v be the vertex of $(G - A)^\star$ corresponding to $K \setminus A$ and v' the vertex of $G^\star$ corresponding to K, for each $x \in V(H)$ such that $f(x) = v$ we add a vertex x' to H' with $\ell'(x') = \ell(x)$ adjacent to the same vertices as x in addition to the vertices of $(f')^{-1}(K \cap A)$, and we set $f'(x') = v'$. We define $L' = (H', f', \ell')$ and we observe that $L' \in \mathcal{P}_r(G)$.

Let $\mathcal{L}' = \{L' : L \in \mathcal{L}\}$. Note that $\mathcal{L}'$ is a system of $\mathcal{P}_r$-overlays of G of treewidth at most $t(r, k) + a$ and thickness $\theta(\mathcal{L}') = \theta(\mathcal{L})$. Hence, the claim of the lemma holds. □

Using the same argument, we can show the following.

Corollary 5.6. *Consider any $\mathcal{P} \in \{\mathcal{A}, \bigstar\}$. Let $\mathcal{G}$ be a $(\mathcal{P}, p, q)$-thin class of graphs for some functions p and q, with t bounding the treewidth of its overlays. Let a, k and r be positive integers, let G be a graph belonging to $\mathcal{G}$ with n vertices, and let A be a set of its vertices of size at most a. In time $O(p(n) + n\,poly(q(n)))$, we can find a system $\mathcal{L}$ of $\mathcal{P}_r$-overlays of G of thickness at most $1 + 1/k$, treewidth at most $t(r, k) + a$, and size at most $poly(q(n))$, such that $\theta_{\mathcal{L}}(v) = 1$ for all $v \in A$.*

Proof. Let $\mathcal{L}'$ be a system of $\mathcal{P}_r$-overlays of G of thickness at most $1 + 1/k$, treewidth at most $t(r,k)$, and size at most $\mathrm{poly}(q(n))$, found in time $O(p(n) + n\,\mathrm{poly}(q(n)))$. By Observation 2.2, we can transform $\mathcal{L}'$ into a system $\mathcal{L}''$ of $\mathcal{P}_r$-overlays of $G - A$ of thickness at most $1 + 1/k$, treewidth at most $t(r,k)$, and size $|\mathcal{L}'|$. The construction from the proof of Lemma 5.5 applied to $\mathcal{L}''$ gives the requested system $\mathcal{L}$ of $\mathcal{P}_r$-overlays of G. □

Let G_0, G_1, ..., G_m be graphs such that for $1 \leq i \leq m$, $G_0 \cap G_i$ is a clique, and for $1 \leq i < j \leq m$, $G_i \cap G_j \subseteq G_0$. We say that the graph $G_0 \cup G_1 \cup \ldots \cup G_m$ is a *star sum of* G_0, ..., G_m, *with center* G_0 *and rays* G_1, ..., G_m. The following lemma gives the main motivation for introducing the notion of ★-thinness.

Lemma 5.7. *For $i \in \{1,2\}$ and some functions p and q, let $\mathcal{G}_i$ be a $(\bigstar, p, q)$-thin class of graphs. Let $\mathcal{G}$ be the class of star sums with center from $\mathcal{G}_1$ and rays from $\mathcal{G}_2$. Suppose furthermore that for each graph $G \in \mathcal{G}$ with n vertices, we can find the center and rays of such a star sum in time $O(p(n) + n\,poly(q(n)))$. Then $\mathcal{G}$ is $(\bigstar, p, q)$-thin.*

Proof. Let t bound the treewidth of overlays of both $\mathcal{G}_1$ and $\mathcal{G}_2$. Let $G \in \mathcal{G}$ be a star sum of $G_0 \in \mathcal{G}_1$ and $G_1, \ldots, G_m \in \mathcal{G}_2$, and let n denote the number of vertices of G. For $i = 1, \ldots, m$, let $A_i = V(G_0 \cap G_i)$. Let $a = 4t(1,1) + 1$. Since A_i induces a clique in G, Lemma 2.3 implies $|A_i| \leq a$.

Let r and k be positive integers. Let $\mathcal{L}_0$ be a system of $\bigstar_r$-overlays of G_0 of thickness at most $1 + \frac{1}{3k}$, treewidth at most $t(r, 3k)$, and size at most $\mathrm{poly}(q(n))$. By Lemma 5.1 and Corollary 5.6, there exist systems $\mathcal{L}_1$, ..., $\mathcal{L}_m$ of $\bigstar_r$-overlays of G_1, ..., G_m all of the same size (at most $poly(q(n))$), thickness at most $1 + \frac{1}{3k}$, and treewidth at most $t(r, 9k) + a$, such that for $i = 1, \ldots, m$, each vertex $v \in A_i$ satisfies $\theta_{\mathcal{L}_i}(v) = 1$. All these systems can be found in total time $O(p(n) + n\,\mathrm{poly}(q(n)))$.

Consider any overlays $L_i = (H_i, f_i, \ell_i) \in \mathcal{L}_i$ for $i = 0, \ldots, m$. For $i \in \{1, \ldots, m\}$, note that $f_i^{-1}(A_i)$ is a clique in H_i and $\ell_i(x) = r$ for all $x \in f_i^{-1}(A_i)$. Let R_i be the set of vertices of $V(G_i^\star) \setminus V(G_i)$ representing cliques that are subsets of A_i, and let $H_i' = H_i - f_i^{-1}(R_i)$. Let v_{A_i} be the vertex of $G^\star$ representing the clique A_i. Let H be the graph obtained from H_0 by, for $i = 1, \ldots, m$ and each vertex $x \in f_0^{-1}(v_{A_i})$, adding a copy $H_{i,x}'$ of the graph H_i' with the clique $f_i^{-1}(A_i)$ identified with the clique on the neighbors of x (so that the functions f_0 and f_i match on the common clique). Let $f : V(H) \to V(G^\star)$ match f_0 on $V(H_0)$ and f_i on each copy $H_{i,x}'$ of H_i' for $i = 1, \ldots, m$. Let $\ell : V(H) \to \{0, 1, \ldots, r\}$ match ℓ_0 on $V(H_0)$, and for $i = 1, \ldots, m$, each vertex $x \in f_0^{-1}(v_{A_i})$, and all vertices $y \in V(H_{i,x}')$, set $\ell(y) = \min(\ell_i(y), \ell_0(x))$. Let us define $L = L(L_0, \ldots, L_m) := (H, f, \ell)$.

Clearly, L is a $\star$-overlay of G. We claim that L is walk-preserving. Indeed, consider any vertex $y \in V(H)$ with $\ell(y) > 0$ and a neighbor v of $f(y)$ in $G^\star$. If $v, f(y) \in V(G_0^\star)$, or if $f(y) \in V(G_i^\star) \setminus V(G_0^\star)$ (and consequently

$v \in V(G_i^\star))$ for some $i \in \{1, \ldots, m\}$, then y has a neighbor z with $\ell(z) \geq \ell(y) - 1$ and $f(z) = v$ by the fact that L_0 or L_i is walk-preserving and by Observation 5.3. Hence, suppose that $f(y) \in V(G_0^\star)$ and $v \in V(G_i^\star) \setminus V(G_0^\star)$ for some $i \in \{1, \ldots, m\}$, and thus $f(y) \in A_i$. Since L_0 is walk-preserving and v_{A_i} is a neighbor of $f(y)$ in $G_0^\star$, there exists a neighbor x of y in H_0 such that $f(x) = v_{A_i}$ and $\ell(x) \geq \ell(y) - 1$. Since L_i is walk-preserving, there exists a neighbor z of y in $H'_{i,x}$ such that $f(z) = v$ and $\ell_i(z) \geq \ell_i(y) - 1 = r - 1$. Consequently, $\ell(z) = \min(\ell_i(z), \ell_0(x)) \geq \ell(y) - 1$.

Furthermore, since $L_0 \in \bigstar_r(G_0)$, for $i = 1, \ldots, m$, there exists $x \in f_0^{-1}(v_{A_i})$ such that $\ell_0(x) = r$. Consequently, since $L_i \in \bigstar_r(G_i)$, for any vertex $v \in V(G_i^\star) \setminus V(G_0^\star)$ there exists a vertex $y \in V(H'_{i,x})$ such that $f(y) = v$ and $\ell(y) = \min(\ell_i(y), \ell_0(x)) = r$. We conclude that $L \in \bigstar_r(G)$.

Let $\mathcal{L}_0 = \{L_0^1, \ldots, L_0^b\}$ and for $i = 1, \ldots, m$, let $\mathcal{L}_i = \{L_i^1, \ldots, L_i^c\}$. For $1 \leq j \leq b$ and $1 \leq l \leq c$, let $L_{j,l} = L(L_0^j, L_1^l, L_2^l, \ldots, L_m^l)$, and let $\mathcal{L} = \{L_{j,l} : 1 \leq j \leq b, 1 \leq l \leq c\}$. Then $\mathcal{L}$ is a system of $\bigstar_r$-overlays of G of size $bc \leq \text{poly}(q(n))$ and treewidth at most $t(r, 9k) + a$. Furthermore, consider any vertex $v \in V(G^\star)$. If $v \in V(G_0^\star)$, then $\theta_{\mathcal{L}}(v) = \theta_{\mathcal{L}_0}(v) < 1 + 1/k$. If $v \in V(G_i^\star) \setminus V(G_0^\star)$ for some $i \in \{1, \ldots, i\}$, then $\theta_{L_{j,l}}(v) = \theta_{L_0^j}(v_{A_i})\theta_{L_i^l}(v)$, and thus

$$\theta_{\mathcal{L}}(v) = \frac{1}{bc} \sum_{j,l} \theta_{L_0^j}(v_{A_i})\theta_{L_i^l}(v) = \theta_{\mathcal{L}_0}(v_{A_i})\theta_{\mathcal{L}_i}(v) \leq \left(1 + \frac{1}{3k}\right)^2 < 1 + 1/k.$$

Hence, $\mathcal{L}$ has thickness at most $1 + 1/k$. □

Next, we consider a generalization of a star sum. Let G be a graph and let $(V_1, \ldots, V_d)$ is its layering. We say that $(V_1, \ldots, V_d)$ is a *shadow-complete layering* if for $i = 1, \ldots, d-1$ and for each connected component C of the graph $G[V_{i+1} \cup V_{i+2} \cup \ldots \cup V_d]$, the neighbors of vertices of C in V_i induce a clique in G.

Lemma 5.8. *Let $\mathcal{G}$ be a $(\bigstar, p, q)$-thin class of graphs for some functions p and q. For a positive integer d, let $\mathcal{G}_d$ be a class of graphs G that have a shadow-complete layering with at most d layers such that each layer induces a graph from $\mathcal{G}$. Furthermore, assume such a layering can be found in time $O(p(n) + n\,\text{poly}(q(n)))$, where $n = |V(G)|$. Then $\mathcal{G}_d$ is $(\bigstar, p, q)$-thin.*

Proof. We prove the claim by induction on d. The case $d = 1$ is trivial, hence assume that $d \geq 2$. Consider a graph $G \in \mathcal{G}_d$ with a shadow-complete layering $(V_1, \ldots, V_d)$ such that each layer induces a subgraph belonging to $\mathcal{G}$. Let $G_0 = G[V_1]$. By Lemma 2.3, each clique in G_0 has size at most $a = 4t(1, 1) + 1$, where t bounds the treewidth of overlays of $\mathcal{G}$.

Let $G' = G[V_2 \cup \ldots \cup V_d]$, and note that $(V_2, \ldots, V_d)$ is a shadow-complete layering of G'; hence, $G' \in \mathcal{G}_{d-1}$, which is a $(\bigstar, p, q)$-thin class by the induction hypothesis. Let $G'_1, \ldots, G'_m$ be the connected components of

G'. For $1 \leq i \leq m$, let A_i be the set of neighbors of G'_i in V_1, and let $G_i = G[V(G'_i) \cup A_i]$. Note that G is a star sum with center G_0 and rays G_1, $\ldots$, G_m. By Lemmas 5.5 and 5.7, we conclude that $\mathcal{G}_d$ is $(\bigstar, p, q)$-thin. □

Using Lemma 5.4, we can lift the previous result to shadow-complete layerings with an unbounded number of layers.

Corollary 5.9. *Let $\mathcal{G}$ be a $(\bigstar, p, q)$-thin class of graphs for some functions p and q. Let $\mathcal{G}'$ be a class of graphs G that have a shadow-complete layering such that each layer induces a graph from $\mathcal{G}$. Furthermore, assume such a layering can be found in time $O(p(n) + n\,\mathrm{poly}(q(n)))$, where $n = |V(G)|$. Then $\mathcal{G}'$ is $(\bigstar, p, q)$-thin.*

Proof. Consider any graph $G \in \mathcal{G}'$ and let $(V_1, \ldots, V_l)$ be its shadow-complete layering with each layer inducing a graph from $\mathcal{G}$. Consider any positive integers d and $i \leq l - d + 1$, and let $G_{i,d} = G[V_i \cup \ldots \cup V_{i+d-1}]$. Observe that $(V_i, \ldots, V_{i+d-1})$ is a shadow-complete layering of $G_{i,d}$, and thus $G_{i,d}$ belongs to the class $\mathcal{G}_d$ from Lemma 5.8. Consequently, $\mathcal{G}'$ is $(\bigstar, p, q)$-thin by Lemma 5.4. □

A *tree decomposition* of a graph G is a pair (T, β), where T is a tree and β assigns to each vertex of T a subset of vertices of G, satisfying the following conditions:

- For each edge $xy \in E(G)$, there exists a vertex $v \in V(T)$ such that $\{u, v\} \subseteq \beta(v)$, and
- for each vertex $x \in V(G)$, the set $\{v \in V(T) : x \in \beta(v)\}$ induces a non-empty connected subtree in T.

The sets $\beta(v)$ for $v \in V(T)$ are called the *bags* of the decomposition. We say that the decomposition is *chordal* if $\beta(v) \cap \beta(v')$ induces a clique in G for every edge $vv' \in E(T)$. A tree decomposition (T', β') of a subgraph G' of G is *contained* in the decomposition (T, β) if for each vertex $v' \in V(T')$, there exists a vertex $v \in V(T)$ such that $\beta(v') \subseteq \beta(v)$. We need the following auxiliary result relating chordal tree decompositions and shadow-complete layerings.

Lemma 5.10 (Dujmović, Morin, and Wood [11])**.** *Let G be a graph that has a chordal tree decomposition (T, β) of adhesion at most a. Then G has a shadow-complete layering such that each layer induces a subgraph of G with a chordal tree decomposition of adhesion at most $a - 1$ contained in (T, β). Furthermore, this layering and the decompositions of each layer can be obtained in linear time given (T, β).*

Together with Corollary 5.9, this enables us to lift $\bigstar$-thinness of a graph class $\mathcal{G}$ to the class of graphs with chordal tree decompositions whose bags belong to $\mathcal{G}$.

Lemma 5.11. *Let $\mathcal{G}$ be a $(\bigstar, p, q)$-thin class of graphs closed on induced subgraphs, for some functions p and q. For a non-negative integer a, let $\mathcal{G}_a$ be a class of graphs G that have a chordal tree decomposition (T, β) of adhesion at most a such that $G[\beta(z)] \in \mathcal{G}$ for all $z \in V(T)$. Furthermore, assume such a decomposition can be found in time $O(p(n) + n\,poly(q(n))$, where $n = |V(G)|$. Then $\mathcal{G}_a$ is $(\bigstar, p, q)$-thin.*

Proof. We prove the claim by induction on a. Consider any graph $G \in \mathcal{G}$ and its tree-decomposition (T, β) as described in the statement of the lemma.

If $a = 0$, then $G[\beta(z)]$ is a component of G for each $z \in V(T)$, and the claim follows from Lemma 5.2. Hence, suppose that $a \geq 1$. Let $(V_1, \ldots, V_l)$ be a shadow-complete layering of G obtained by Lemma 5.10. Consider any $i \in \{1, \ldots, l\}$, and let (T_i, β_i) be a chordal tree decomposition of adhesion at most $a - 1$ of $G[V_i]$ contained in (T, β). Since $\mathcal{G}$ is closed on induced subgraphs, $G[\beta_i(z)] \in \mathcal{G}$ for all $z \in V(T_i)$. Hence $G[V_i]$ belongs to $\mathcal{G}_{a-1}$, which is $(\bigstar, p, q)$-thin by the induction hypothesis. Consequently, $\mathcal{G}_a$ is $(\bigstar, p, q)$-thin by Corollary 5.9. □

Let us restate the previous result in the terms of arbitrary (not necessarily chordal) tree decompositions. Let (T, β) be a tree decomposition of a graph G. The *torso* of a vertex $z \in V(T)$ is obtained from $G[\beta(z)]$ by adding cliques on $\beta(z) \cap \beta(z')$ for all $zz' \in E(T)$.

Corollary 5.12. *Let $\mathcal{G}$ be a $(\bigstar, p, q)$-thin class of graphs closed on induced subgraphs, for some functions p and q. Let $\mathcal{G}'$ be a class of graphs G that have a tree decomposition (T, β) such that the torso of each vertex of T belongs to $\mathcal{G}$. Furthermore, assume such a decomposition can be found in time $O(p(n) + n\,poly(q(n))$, where $n = |V(G)|$. Then $\mathcal{G}'$ is $(\bigstar, p, q)$-thin.*

Proof. Let t bound the treewidth of overlays of $\mathcal{G}$. Let G be a graph from $\mathcal{G}'$ and let (T, β) be its tree decomposition whose torsos belong to $\mathcal{G}$. By Lemma 2.3, each graph $\mathcal{G}$ has clique number at most $a = 4t(1, 1) + 1$, and hence the decomposition (T, β) has adhesion at most a. Let G' be the graph obtained from G by adding cliques on sets $\beta(z) \cap \beta(z')$ for all $zz' \in E(T)$. Then (T, β) is a chordal tree decomposition of G' and $G'[\beta(z)] \in \mathcal{G}$ for each $z \in V(T)$. Hence, G' belongs to the $(\bigstar, p, q)$-thin class $\mathcal{G}_a$ from Lemma 5.11. Observation 2.2 implies that $\mathcal{G}'$ is $(\bigstar, p, q)$-thin. □

Note that equivalently, $\mathcal{G}'$ is a class of graphs obtained from those in $\mathcal{G}$ by clique-sums.

5.2 Proper minor-closed classes

We are now ready to show that proper minor-closed classes are thin. Each graph G has a trivial r-neighborhood overlay (G, id, r). Hence, we get the following.

Observation 5.13. *Every class of graphs of bounded treewidth is* $(\mathcal{S}, n, 1)$*-thin.*

Dujmović, Morin, and Wood [11] introduced the following notion. Let c be a positive integer. A graph G has *layered treewidth at most* c if G has a tree decomposition (T, β) and a layering $(V_1, \ldots, V_l)$ such that $|V_i \cap \beta(z)| \leq c$ for $1 \leq i \leq l$ and each vertex $z \in V(T)$. Note that union of any d layers induces a subgraph of G of treewidth at most cd. Observation 5.13 together with Lemma 5.4 implies the following.

Corollary 5.14. *Let* c *be a positive integer and let* p *be a positive non-decreasing polynomial. Suppose that* $\mathcal{G}$ *is a class of graphs of layered treewidth at most* c*. If the corresponding tree decomposition and layering of an* n*-vertex graph from* $\mathcal{G}$ *can be found in time* $O(p(n))$*, then* $\mathcal{G}$ *is* $(\mathcal{S}, p, 1)$*-thin.*

Let us state a corollary of Robertson-Seymour structure theorem, see Grohe [17] and Dujmović et al. [11] for more details.

Theorem 5.15. *For every proper minor-closed class* $\mathcal{G}$*, there exist positive integers* a *and* c *as follows. Every graph* $G \in \mathcal{G}$ *has a tree decomposition* (T, β) *and a system* $\{A_z \subseteq \beta(z) : z \in V(T)\}$ *of sets such that for every* $z \in V(T)$*, the set* A_z *has size at most* a *and letting* F *denote the torso of* z*, the graph* $F - A_z$ *has layered treewidth at most* c*. Furthermore, the tree decomposition of* G *and the system* $\{A_z : z \in V(T)\}$ *can be found in time* $O(|V(G)|^3)$*, and the decomposition and layering of* $F - A$ *for each vertex can be found in linear time.*

Corollary 5.14, Lemmas 4.2 and 5.5, and Corollary 5.12 imply the following.

Corollary 5.16. *Every proper minor-closed class* $\mathcal{G}$ *is* $(\bigstar, n^3, 1)$*-thin.*

5.3 Strongly sublinear separators

A *separation* of a graph G is a pair (A, B) of edge-disjoint subgraphs of G such that $A \cup B = G$, and the *size* of the separation is $|V(A) \cap V(B)|$. Observe that G has no edge with one end with $V(A) \setminus V(B)$ and the other end in $V(B) \setminus V(A)$, and thus the set $V(A) \cap V(B)$ separates $V(A) \setminus V(B)$ from $V(B) \setminus V(A)$ in G. A separation (A, B) is *balanced* if $|V(A) \setminus V(B)| \leq 2|V(G)|/3$ and $|V(B) \setminus V(A)| \leq 2|V(G)|/3$. Note that $(G, G - E(G))$ is a balanced separation. For a graph class $\mathcal{G}$, let $s_{\mathcal{G}}(n)$ denote the smallest nonnegative integer such that every graph in $\mathcal{G}$ with at most n vertices has a balanced separation of size at most $s_{\mathcal{G}}(n)$. We say that $\mathcal{G}$ has *sublinear separators* if $\lim_{n\to\infty} \frac{s_{\mathcal{G}}(n)}{n} = 0$, and that $\mathcal{G}$ has *strongly sublinear separators* if there exist constants $c \geq 1$ and $0 \leq \delta < 1$ such that $s_{\mathcal{G}}(n) \leq cn^{\delta}$ for every $n \geq 0$.

We now aim to show that subgraph-closed graph classes with strongly sublinear separators and bounded maximum degree are $\mathcal{S}$-thin. We start by obtaining layerings of such graphs from their tree decompositions. We use the following fact about tree decompositions of graphs of bounded maximum degree.

Lemma 5.17 (Dvořák [12])**.** *Let G be a graph of maximum degree at most Δ. If G has treewidth at most t, then G has a rooted tree decomposition (T, β) with bags of size at most $12(t+1)(\Delta+1)$ such that every vertex of T has at most two sons and for each $v \in V(G)$, the subtree $T[\{u : v \in \beta(u)\}]$ has depth at most $1 + 4\log(\Delta+1)$.*

The proof of this lemma gives an algorithm to find such a decomposition in time $O(t\Delta|V(G)|^2)$, assuming that a tree decomposition of G of width at most t is given.

Let $\mathcal{G}$ be a class of graphs closed on induced subgraphs. If $\mathcal{G}$ has strongly sublinear separators, then each graph $G \in \mathcal{G}$ has sublinear treewidth; a tree decomposition of sublinear width can be constructed recursively as follows. Find a balanced separator (A, B) of order at most $s_{\mathcal{G}}(|V(G)|)$, find tree decompositions of $A - V(B)$ and $B - V(A)$ inductively, and obtain a tree decomposition of G by joining the two by an arbitrary edge and adding $V(A \cap B)$ to all the bags. Hence, we obtain the following.

Observation 5.18. *If s is a positive integer and every induced subgraph of an n-vertex graph G has a balanced separator of order at most s, then G has a tree decomposition of width at most $s\lceil\log_{3/2} n\rceil$. Furthermore, if a balanced separator of order at most s in each induced subgraph G' of G can be found in time $O(p(|V(G')|))$ for a positive non-decreasing polynomial p of degree greater than 1, then such a decomposition can be constructed in time $O(p(n))$.*

Dvořák and Norin [13] proved that every subgraph-closed class of graphs with strongly sublinear separators has polynomial expansion. Using algorithm of Plotkin et al. [23], a balanced separator of strongly sublinear size can be obtained in (less than) quadratic time (see Dvořák and Norin [13, Corollary 2] for details). Hence, we obtain the following.

Corollary 5.19. *Let $\mathcal{G}$ be a subgraph-closed class of graphs with strongly sublinear separators. Then there exists a nonnegative real number $\delta < 1$ such that for an n-vertex graph $G \in \mathcal{G}$, a tree decomposition of width $O(n^\delta)$ can be found in time $O(n^2)$.*

We now combine the results to obtain layerings where unions of a small number of consecutive layers induce subgraphs whose all connected components are small.

Corollary 5.20. *Let $\mathcal{G}$ be a subgraph-closed class of graphs with strongly sublinear separators and bounded maximum degree. There exists a positive integer c and a non-negative real number $\delta < 1$ as follows. Every n-vertex graph $G \in \mathcal{G}$ has a layering such that the union of any b consecutive layers induces a subgraph whose components have size at most $c^b n^\delta$. Furthermore, such a layering can be found in time $O(n^3)$.*

Proof. By Corollary 5.19 and Lemma 5.17, there exist positive integers a and c' and a non-negative real number $\delta < 1$ such that for any n-vertex graph $G \in \mathcal{G}$, we can in time $O(n^3)$ find a tree decomposition (T, β) of G of width at most $c'n^\delta$, where T is a rooted tree in which each vertex has at most two sons and for each $v \in V(G)$, the subtree $T[\{u : v \in \beta(u)\}]$ has depth at most a. Let $U_0 = \emptyset$. For $i \geq 1$, let U_i be the union of $\beta(z)$ for all vertices $z \in V(T)$ whose distance from the root of T is at least $(i-1)a$ and less than ia, and let $V_i = U_i \setminus U_{i-1}$. Let d be maximum integer such that V_d is non-empty. Observe that $(V_1, \ldots, V_d)$ is a layering of G.

For $1 \leq i \leq d$ and $1 \leq b \leq d - i + 1$, let $V_{i,b} = V_i \cup V_{i+1} \cup \ldots \cup V_{i+b-1}$ and let $\beta_{i,b}$ be the function defined by $\beta_{i,b}(z) = \beta(z) \cap V_{i,b}$. Let $X_{i,b}$ be the set of vertices $z \in V(T)$ such that $\beta_{i,b}(z) \neq \emptyset$. Note that $(T, \beta_{i,b})$ is a tree decomposition of $G[V_{i,b}]$, and that each component of $T[X_{i,b}]$ is a rooted tree of depth at most $a(b+2)$. Consequently, each component of $G[V_{i,b}]$ has size at most $2^{a(b+2)}c'n^\delta$, and thus Corollary 5.20 holds with $c = c'2^{a+2}$. □

We now proceed similarly to Lemma 5.4.

Lemma 5.21. *Let $\mathcal{G}$ be a subgraph-closed class of graphs with strongly sublinear separators and bounded maximum degree. Let c and δ be the corresponding constants from Corollary 5.20. Let r, k, and t be positive integers, and let α be a positive real number. Let $n_1 = t$, $s_1 = 1$, and $\theta_1 = 1$, and for $i \geq 2$, let $s_i = \lceil 4\alpha^{-1}(1+\alpha)^i rk \rceil$, $n_i = \left(\frac{n_{i-1}}{c^{s_i+2r}}\right)^{1/\delta}$, and $\theta_i = \theta_{i-1} + 4r/s_i$. For an n-vertex graph $G \in \mathcal{G}$ and every positive integer i such that each component of G has at most n_i vertices, we can in time $O(n^3)$ find a system of $\mathcal{S}_r$-overlays of G of thickness at most θ_i, treewidth at most t, and size exactly $s_1 \cdot \ldots \cdot s_i$. Furthermore, $\theta_i \leq 1 + 1/k$.*

Proof. We prove the claim by induction on i. The case $i = 1$ is trivial by Observation 5.13, hence assume that $s \geq 2$. Consider an n-vertex graph $G \in \mathcal{G}$ such that each component of G has at most n_i vertices. Suppose that G has $m \geq 2$ connected components and that we can find systems $\mathcal{L}_1$, …, $\mathcal{L}_m$ of $\mathcal{S}_r$-overlays of the components of G with the described properties; then $\bigcirc_{j=1}^m \mathcal{L}_j$ is a system of $\mathcal{S}_r$-overlays of G of thickness at most θ_i, treewidth at most t, and size $s_1 \cdot \ldots \cdot s_i$. Hence, we can without loss of generality assume that G is connected, and consequently $n \leq n_i$.

Let $(V_1, \ldots, V_d)$ be the layering of G obtained by Corollary 5.20. For integers j such that $j \leq 0$ or $j > d$, let $V_j = \emptyset$. For an integer j, let G_j be

the subgraph of G induced by $V_{j-r} \cup \ldots \cup V_{j+s_i+r-1}$. Each component of G_j has size at most $c^{s_i+2r} n_i^\delta = n_{i-1}$. For each j such that $V(G_j) \neq \emptyset$, let $\mathcal{L}_j$ be the system of $\mathcal{S}_r$-overlays of G_j of thickness at most θ_{i-1}, treewidth at most t, and size exactly $s_1 \cdot \ldots \cdot s_{i-1}$, obtained by the induction hypothesis. Let $\mathcal{L}'_j$ be the system of overlays of G obtained from $\mathcal{L}_j$ as in the proof of Lemma 5.4. For $0 \leq b \leq s_i - 1$, let

$$\mathcal{L}^b = \bigcirc_{\substack{j \bmod \Delta = b, \\ V(G_j) \neq \emptyset}} \mathcal{L}'_j,$$

and let $\mathcal{L} = \bigcup_{b=0}^{s_i-1} \mathcal{L}^b$. As in the proof of Lemma 5.4, we argue that $\mathcal{L}$ is a system of $\mathcal{S}_r$-overlays of G of thickness at most $\theta_{i-1}(1 + 2r/s_i) \leq \theta_{i-1} + 4r/s_i = \theta_i$, treewidth at most t and size $s_1 \cdot \ldots \cdot s_i$.

The last inequality uses the fact that

$$\begin{aligned}\theta_i &\leq 1 + 4r \sum_{j=1}^{\infty} \frac{1}{s_j} \leq 1 + \frac{\alpha}{k} \sum_{j=1}^{\infty} (1+\alpha)^{-j} \\ &= 1 + \frac{\alpha}{k} \cdot \frac{1}{\alpha} = 1 + 1/k \leq 2.\end{aligned}$$

□

The thinness of subgraph-closed class of graphs with strongly sublinear separators and bounded maximum degree now follows simply by a careful choice of parameters. Note that $e^{(\log \log n)^2} = O(e^{\varepsilon \log n}) = O(n^\varepsilon)$ for every $\varepsilon > 0$, and thus the function $(\log \log n)^2$ is subpolynomial.

Corollary 5.22. *Every subgraph-closed class $\mathcal{G}$ of graphs with strongly sublinear separators and bounded maximum degree is $(\mathcal{S}, n^3, e^{(\log \log n)^2})$-thin.*

Proof. Let c and δ be the corresponding constants from Corollary 5.20, and let $\varepsilon = \min(1/\delta - 1, 1/2)$. Let r and k be positive integers, let $c' = c^{2r}$, $\alpha = \varepsilon(1-\varepsilon)/2$, and $a = c^{8\alpha^{-1}rk}$. Let $t = t(r,k) = \max\left(3, \left(c' a^{(1+\alpha)^2}\right)^{2/\varepsilon}\right)$. Let s_i and n_i for any positive integer i be as in the statement of Lemma 5.21. Note that $s_i \leq 8\alpha^{-1}rk(1+\alpha)^i$ for all $i \geq 1$. We have

$$n_i \geq \left(\frac{n_{i-1}}{c' a^{(1+\alpha)^i}}\right)^{1+\varepsilon}.$$

By induction on i, we will show that

$$c' a^{(1+\alpha)^{i+1}} \leq n_i^{\varepsilon/2}$$

holds for every $i \geq 1$. Indeed, the choice of $t = n_1$ ensures that the base case $i = 1$ is true. For $i \geq 2$, we have

$$\begin{aligned} n_i^{\varepsilon/2} &\geq \left(\frac{n_{i-1}}{c' a^{(1+\alpha)^i}}\right)^{(1+\varepsilon)\varepsilon/2} \geq \left(c' a^{(1+\alpha)^i}\right)^{(2/\varepsilon - 1)(1+\varepsilon)\varepsilon/2} \\ &= \left(c' a^{(1+\alpha)^i}\right)^{(1-\varepsilon/2)(1+\varepsilon)} = \left(c' a^{(1+\alpha)^i}\right)^{1+\alpha} \\ &\geq c' a^{(1+\alpha)^{i+1}}, \end{aligned}$$

where the second inequality holds by the induction hypothesis. We conclude that for $i \geq 1$, we have

$$n_i \geq n_{i-1}^{1+\alpha} \geq \ldots \geq n_1^{(1+\alpha)^{i-1}} \geq 3^{(1+\alpha)^{i-1}}.$$

Consider an n-vertex graph $G \in \mathcal{G}$, and let $i = \left\lceil \frac{1}{\log(1+\alpha)} \log\log n \right\rceil + 1$, so that $n_i \geq n$. By Lemma 5.21, we can in time $O(n^3)$ find a system of $\mathcal{S}_r$-overlays of G of thickness at most $1 + 1/k$, treewidth at most $t(r,k)$, and size at most $s_1 \cdot \ldots \cdot s_i \leq s_i^i = \text{poly}(e^{(\log\log n)^2})$. □

References

[1] Baker, B. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM (JACM) 41*, 1 (1994), 153–180.

[2] Berger, E., Dvořák, Z., and Norin, S. Treewidth of grid subsets. *Combinatorica* (2017). accepted.

[3] Cabello, S., and Gajser, D. Simple PTAS's for families of graphs excluding a minor. *Discrete Applied Mathematics 189* (2015), 41–48.

[4] Courcelle, B. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and computation 85*, 1 (1990), 12–75.

[5] Dawar, A., Grohe, M., Kreutzer, S., and Schweikardt, N. Approximation schemes for first-order definable optimisation problems. In *21st Annual IEEE Symposium on Logic in Computer Science (LICS'06)* (2006), IEEE, pp. 411–420.

[6] Demaine, E., Hajiaghayi, M., and Kawarabayashi, K. Contraction decomposition in H-minor-free graphs and algorithmic applications. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing* (2011), STOC '11, ACM, pp. 441–450.

[7] Demaine, E. D., and Hajiaghayi, M. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms* (2004), Society for Industrial and Applied Mathematics, pp. 840–849.

[8] Demaine, E. D., and Hajiaghayi, M. Bidimensionality: new connections between FPT algorithms and PTASs. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms* (2005), Society for Industrial and Applied Mathematics, pp. 590–601.

[9] Demaine, E. D., Hajiaghayi, M. T., and Kawarabayashi, K.-i. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)* (2005), IEEE, pp. 637–646.

[10] DeVos, M., Ding, G., Oporowski, B., Sanders, D., Reed, B., Seymour, P., and Vertigan, D. Excluding any graph as a minor allows a low tree-width 2-coloring. *J. Comb. Theory, Ser. B 91* (2004), 25–41.

[11] Dujmović, V., Morin, P., and Wood, D. R. Layered separators in minor-closed families with applications. *arXiv e-prints 1306.1595* (2013).

[12] Dvořák, Z. Sublinear separators, fragility and subexponential expansion. *European Journal of Combinatorics 52, Part A* (2016), 103–119.

[13] Dvořák, Z., and Norin, S. Strongly sublinear separators and polynomial expansion. *SIAM Journal on Discrete Mathematics 30* (2016), 1095–1101.

[14] Eppstein, D. Diameter and treewidth in minor-closed graph families. *Algorithmica 27* (2000), 275–291.

[15] Fomin, F. V., Lokshtanov, D., Raman, V., and Saurabh, S. Bidimensionality and EPTAS. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms* (2011), SIAM, pp. 748–759.

[16] Grigoriev, A., and Bodlaender, H. L. Algorithms for graphs embeddable with few crossings per edge. *Algorithmica 49*, 1 (2007), 1–11.

[17] Grohe, M. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica 23*, 4 (2003), 613–632.

[18] Grohe, M., Kreutzer, S., and Siebertz, S. Deciding first-order properties of nowhere dense graphs. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (2014), ACM, pp. 89–98.

[19] Har-Peled, S., and Quanrud, K. Approximation algorithms for polynomial-expansion and low-density graphs. In *Algorithms-ESA 2015*. Springer, 2015, pp. 717–728.

[20] Hunt, H. B., Marathe, M. V., Radhakrishnan, V., Ravi, S. S., Rosenkrantz, D. J., and Stearns, R. E. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms 26*, 2 (1998), 238–274.

[21] Nešetřil, J., and Ossona de Mendez, P. Grad and classes with bounded expansion II. Algorithmic aspects. *European J. Combin. 29* (2008), 777–791.

[22] Papadimitriou, C., and Yannakakis, M. Optimization, approximation, and complexity classes. In *Proceedings of the twentieth annual ACM symposium on Theory of computing* (1988), ACM, pp. 229–234.

[23] Plotkin, S., Rao, S., and Smith, W. D. Shallow excluded minors and improved graph decompositions. In *Proceedings of the fifth annual ACM-SIAM symposium on Discrete algorithms* (1994), Society for Industrial and Applied Mathematics, pp. 462–470.